UNITED STATES PATENT AND TRADEMARK OFFICE

TH

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

| APPLICATION NO. | FILING DATE | FIRST NAMED INVENTOR | ATTORNEY DOCKET NO. | CONFIRMATION NO. |
|---|---|---|---|---|
| 10/614,338 | 07/08/2003 | Farshid Nowshadi | 060707-1280 | 7873 |

24504 7590 03/27/2007

THOMAS, KAYDEN, HORSTEMEYER & RISLEY, LLP
100 GALLERIA PARKWAY, NW
STE 1750
ATLANTA, GA 30339-5948

| EXAMINER | |
|---|---|
| RICHMOND, LEAH L | |
| ART UNIT | PAPER NUMBER |
| 2609 | |

| SHORTENED STATUTORY PERIOD OF RESPONSE | MAIL DATE | DELIVERY MODE |
|---|---|---|
| 3 MONTHS | 03/27/2007 | PAPER |

**Please find below and/or attached an Office communication concerning this application or proceeding.**

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

| Office Action Summary | Application No. | Applicant(s) | |
|---|---|---|---|
| | 10/614,338 | NOWSHADI, FARSHID | |
| | Examiner | Art Unit | |
| | Leah L. Richmond | 2609 | |

*-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --*

**Period for Reply**

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

**Status**

1)☒ Responsive to communication(s) filed on *08 July 2003*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

**Disposition of Claims**

4)☒ Claim(s) *1-26* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-26* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

**Application Papers**

9)☒ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *21 October 2003* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

**Priority under 35 U.S.C. § 119**

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

**Attachment(s)**

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *12/09/2003; 04/19/2004*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

# DETAILED ACTION

## *Objections to Specification*

Page 9, lines 25 – 26, of the specification states: "Figure 7 is an exemplary diagram illustrating a system for allocating a shared resource in accordance with an embodiment of the present invention. A system 700 for ..." There is no label "700" in Fig. 7.

Page 10, line 8, of the specification states: "A Grant Access Module 726 may granting access to ..." referring to Fig. 7. In Fig. 7, the Grant Access Module is labeled 728.

Appropriate correction is required.

## *Claim Rejections - 35 USC § 102*

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office Action:

> A person shall be entitled to a patent unless -- (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

**Claims 1** and **14** are rejected under 35 U.S.C. 102(b) as being anticipated by **Iverson et al. (U.S. Patent # 6052379)**.

Consider **claim 1**, Iverson et al. clearly show and disclose a method for allocating a shared resource among a plurality of devices, comprising associating a bucket to each device, assigning a rate at which each bucket fills with credits, assigning a rate at

which each bucket drains credits, comparing the fill level of each bucket to determine which bucket has the most credits, and prioritizing access to the shared resource according to the credit level of the buckets (Fig. 10 and Col. 2, lines 21 – 27: "The invention provides a mechanism for fair, low latency packet-based bandwidth delivery in a TDM system using a priority scheme based on a "leaky bucket" mechanism. Bandwidth delivery to the TDM system is controlled by the priority assigned to user traffic. The priority scheme uses the "water level" in the leaky bucket priority scheme to fairly assign priority within a high or low priority band." and Col. 2, lines 40 - 53: "The priority scheme maintains two buckets to track the current bandwidth delivery rate and unused committed bandwidth. The size of the first bucket is dimensioned as twice the configured committed burst capacity ($B_c$) for a given User connection. The water level in the first bucket represents the current bandwidth delivery rate for the User. If the water level is above the midpoint and is rising, the User is sending traffic below their CIR and is accumulating credit for the unused allocated bandwidth. If the water level is above the midpoint but is falling, the User is sending traffic above their CIR but has not exceeded the committed burst capacity configured. If the water level is below the bucket midpoint, the User has used up their committed burst capacity." And Col. 3, lines 1-10: "If one considers the midpoint of the first bucket to be the origin, water levels above the midpoint of the first bucket are considered positive and those below are considered negative with the midpoint being equal to zero. When the water level in the first bucket is above the zero, the priority is derived from the water level in the first bucket relative to the range from zero to the top of the bucket establishing a high priority

band. The actual priority value is determined by a discrete mapping of the water level onto a range of high priority values in the high priority band.").

Consider **claim 14**, Iverson et al. clearly show and disclose a system for allocating a shared resource among a plurality of devices, comprising an association module for assigning a bucket to each device, a fill rate module for assigning a rate at which each bucket fills with credits, a drain rate module for assigning a rate at which each bucket drains credits, a grant determination module for comparing the fill level of each bucket to determine which bucket has the most credits, and a grant access module for prioritizing access to the shared resource according to the credit level of the buckets (Fig. 3 and Col. 5, lines 53 – 55: "Fig. 3 is a block diagram of a system for communicating packetized data over a channel bank in accordance with the invention." and Fig. 10 and Col. 2, lines 21 – 27: "The invention provides a mechanism for fair, low latency packet-based bandwidth delivery in a TDM system using a priority scheme based on a "leaky bucket" mechanism. Bandwidth delivery to the TDM system is controlled by the priority assigned to user traffic. The priority scheme uses the "water level" in the leaky bucket priority scheme to fairly assign priority within a high or low priority band." and Col. 2, lines 40 - 53: "The priority scheme maintains two buckets to track the current bandwidth delivery rate and unused committed bandwidth. The size of the first bucket is dimensioned as twice the configured committed burst capacity ($B_c$) for a given User connection. The water level in the first bucket represents the current bandwidth delivery rate for the User. If the water level is above the midpoint and is rising, the User is sending traffic below their CIR and is accumulating credit for the

unused allocated bandwidth. If the water level is above the midpoint but is falling, the User is sending traffic above their CIR but has not exceeded the committed burst capacity configured. If the water level is below the bucket midpoint, the User has used up their committed burst capacity." And Col. 3, lines 1-10: "If one considers the midpoint of the first bucket to be the origin, water levels above the midpoint of the first bucket are considered positive and those below are considered negative with the midpoint being equal to zero. When the water level in the first bucket is above the zero, the priority is derived from the water level in the first bucket relative to the range from zero to the top of the bucket establishing a high priority band. The actual priority value is determined by a discrete mapping of the water level onto a range of high priority values in the high priority band.").

***Claim Rejections - 35 USC § 103***

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

The factual inquiries set forth in *Graham* **v.** *John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.

4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

**Claims 2, 3, 15** and **16** are rejected under 35 U.S.C. 103(a) as being unpatentable over **Iverson et al. (U.S. Patent # 6052379)** in view of **Rundberg ("Memory Bandwidth Explained")**.

Consider **claims 2** and **15**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except for a shared resource that is a memory bandwidth. However, Rundberg clearly shows and discloses a shared resource that is a memory bandwidth (Introduction, Paragraph 2: "New memory technologies like DRDRAM (Rambus) and DDR-SDRAM are also becoming common. These new memory technologies have in common that they provide increased peak memory bandwidth by enhancing the bus between the memory and the microprocessor."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate a shared resource that is a memory bandwidth, as taught by Rundberg, in the method and system for allocating a shared resource among a plurality of devices, as in Iverson et al. for the purpose of allocating memory bandwidth among a plurality of devices.

Consider **claims 3** and **16**, and as applied to claims 2 and 15, Iverson et al. disclose the claimed invention except for a memory bandwidth associated with one or more of DRAM, SDRAM, SRAM and EPROM. However, Rundberg clearly shows and discloses memory bandwidth that is associated with DRAM (Introduction, paragraph 2: "New memory technologies like DRDRAM (Rambus) and DDR-SDRAM are also becoming common." and "The latency is not so much an issue of the memory interface

as it is the memory cell itself, and since both these two new memories use DRAM, the latency is not improved."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate memory bandwidth that is associated with one or more of DRAM, SDRAM, SRAM and EPROM, as taught by Rundberg, in the method and system for allocating a shared resource among a plurality of devices, as in Iverson et al. for the purpose of allocating memory bandwidth that is associated with new memory technologies.

**Claims 4, 5, 17** and **18** are rejected under 35 U.S.C. 103(a) as being unpatentable over **Iverson et al. (U.S. Patent # 6052379)** in view of **Minerd ("Time Division Multiplexed Bus for Machine Control")**.

Consider **claims 4** and **17**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except for a shared resource comprising a bus connected to at least one peripheral device including one or more of TDM, UART, USB, and PCI. However, Minerd clearly shows and discloses a shared resource comprising a bus connected to at least one peripheral device including TDM (MUX BUS TDM Signals Section: "The MUX BUS consists of three time division multiplexed signals (SID, SOD, and CLK) that are used by the base controller to manipulate machine I/O and two supply lines to power the modules (see Fig. 4)."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate a bus including at least one peripheral device including TDM as taught by

Minerd in the method and system as in Iverson et al. for the purpose of allocating a shared resource that is a bus among a plurality of devices.

Consider **claims 5** and **18**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except that the plurality of devices comprise processing units. However, Minerd clearly shows and discloses a method and system wherein the plurality of devices comprise processing units (Fig. 5 and IIOC Chip Section: "The key to understanding the capabilities of the IIOC module is to understand the IIOC chip (see Fig. 5). The IIOC chip is a custom IC manufactured in a BCD process. The IC circuitry is able to interface to the MUX BUS to extract the desired output machine load state and to return the state of its input machine loads." and Fig. 2 and IIOC Architecture Section: "The MUX BUS is shared by up to 16 modules."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate processing units as the plurality of devices as taught by Minerd in the method and system as in Iverson et al. for the purpose of allocating a shared resource among a plurality of processing units.

**Claims 6** and **19** are rejected under 35 U.S.C. 103(a) as being unpatentable over **Iverson et al. (U.S. Patent # 6052379)** in view of **Shah ("Understanding Network Processors")**.

Consider **claims 6** and **19**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except that the plurality of devices comprise at least a combination of a DMA controller, a network processor and a protocol processor.

However, Shah clearly shows and discloses a plurality of devices that comprise at least a combination of a DMA controller, a network processor and a protocol processor (page 43, Fig. 27 and page 42, Section 3.13 IBM (PowerNP), Architecture: "IBM's network processor consists of the Embedded Processor Complex (EPC), special frame processing hardware, and peripheral interfaces. The EPC has a PowerPC core and 16 programmable protocol processors (that make up the Embedded Processor Complex) ... The seven co-processors execute the following functions: Data store: interfaces frame buffer to provide DMA capability ..."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate a plurality of devices comprising at least a combination of a DMA controller, a network processor and a protocol processor, as taught by Shah, in the method and system for allocating a shared resource among a plurality of devices for the purpose of allocating a shared resource among particular types of devices.

**Claims 7, 11, 20** and **24** are rejected under 35 U.S.C. 103(a) as being unpatentable over **Iverson et al. (U.S. Patent # 6052379)** in view of **Parekh et al. ("A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single-Node Case").**

Consider **claims 7** and **20**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except that each fill rate is different, each fill rate indicating access priority assigned to the associated device. However, Parekh et al. clearly show and disclose a method and system wherein each fill rate is different, each

fill rate indicating access priority assigned to the associated device (Section II. GPS Multiplexing: "A good scheme should allow the network to treat users differently, in accordance with their desired quality of service." and "session *i* is guaranteed a rate of $g_i = (\varphi_i)r / (\Sigma_j \varphi_j)$." and "By varying the $\varphi_i$'s, we have the flexibility of treating the sessions in a variety of different ways. For example, when all $\varphi_i$'s are equal, the system reduces to uniform processor sharing. As long as the combined average rate of the sessions is less than *r*, any assignment of positive $\varphi_i$'s yields a stable system." and "When $2\varphi_1 = \varphi_2$ and both sessions are backlogged, session 1 is served at rate 1/3 and session 2 at rate 2/3."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate the method wherein each fill rate is different as taught by Parekh et al. in the method and system for allocating a shared resource among a plurality of devices as in Iverson et al. for the purpose of giving the devices different levels of priority access to the shared resource.

Consider **claims 11** and **24**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except for dynamically adjusting one or more of the fill rate and drain rate associated with one or more buckets for load balancing one or more of the plurality of devices. However, Parekh et al. clearly show and disclose a method and system further comprising the step of dynamically adjusting one or more of the fill rate and drain rate associated with one or more buckets for load balancing (Section II. GPS Multiplexing: "session *i* is guaranteed a rate of $g_i = (\varphi_i)r / (\Sigma_j \varphi_j)$." and "By varying the $\varphi_i$'s, we have the flexibility of treating the sessions in a variety of different ways. For example, when all $\varphi_i$'s are equal, the system reduces to uniform processor sharing. As

long as the combined average rate of the sessions is less than *r*, any assignment of positive $\varphi_i$'s yields a stable system. For example, a high-bandwidth delay-insensitive session *i* can be assigned $g_i$ much less than its average rate, thus allowing for better treatment of the other sessions."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate dynamically adjusting one or more of the fill rate and drain rate for load balancing as taught by Parekh et al. in the method and system for allocating a shared resource among a plurality of devices as in Iverson et al. for the purpose of using the shared resource efficiently.

**Claims 8** and **21** are rejected under 35 U.S.C. 103(a) as being unpatentable over **Iverson et al. (U.S. Patent # 6052379)** in view of **Ng et al. (U.S. Patent # 6157978).**

Consider **claims 8** and **21**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except that each drain rate is different, each drain rate indicating access priority assigned to the associated device. However, Ng et al. clearly show and disclose a method and system wherein each drain rate is different, each drain rate indicating access priority assigned to the associated device (Col. 9, lines 10 - 30: "Once timer 40 has been triggered, a priority arbiter determines which of the non-real-time agents requesting memory access has the highest priority. This highest-priority non-real-time agent is serviced first. In Fig. 6, agent F is serviced first. Once NRT agent F has completed, such as by filling its FIFO, the next-highest-priority NRT agent is selected. In this example, NRT agent G is given access. Agent G has no FIFO, but

merely performs some number of reads or writes of the shared multimedia memory. The next NRT agent is priority order is NRT agent H, which begins to fill its FIFO. Before NRT agent H can finish, timer 40 signals the NRT timeout, and NRT agent H must stop, leaving its FIFO only partially filled."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate varying drain rates, indicating access priority assigned to the associated device, as taught by Ng et al. in the method and system for allocating a shared resource among a plurality of devices as in Iverson et al. for the purpose of prioritizing device access time to the shared resource.

**Claims 12, 13, 25** and **26** are rejected under 35 U.S.C. 103(a) as being unpatentable over **Iverson et al. (U.S. Patent # 6052379)** in view of **Boivie et al. (U.S. Patent # 6842783).**

Consider **claims 12** and **25**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except determining an amount of bandwidth each device has used. However, Boivie et al. clearly show and disclose the step of determining an amount of bandwidth each device has used (Fig. 2 and Col. 6, lines 5 – 14: "The path 160 carries the outbound traffic (i.e., responses to requests). Hence, monitoring the output on the outgoing data path 160 can be monitored and it can be observed how much output is being generated. If the amount of bandwidth being used by a particular customer exceeds the number of bandwidth per the SLA, then feedback can be generated to reduce the number of inbound requests being accepted. Thus, the

present invention provides a feedback loop, so as to provide monitored data 260 indicating the bandwidth on output path 160 being used by each of the Web sites hosted on the Web server cluster 100."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate the step of determining how much bandwidth each device has used, as taught by Boivie et al., in the method and system for allocating a shared resource among a plurality of devices, as in Iverson et al., for the purpose of providing feedback to use in enforcing minimum and maximum bandwidth limits for each device.

Consider **claims 13** and **26**, and as applied to claims 12 and 25, Iverson et al. disclose the claimed invention except charging an entity according to the amount of bandwidth used. However, Boivie et al. clearly show and disclose charging an entity according to the amount of bandwidth used (Fig. 4A and Col. 7, line 53: "Step 406 guarantees the minimum bandwidth SLA." and Col. 7, line 66 – Col. 8, line 4: "The selection of a class is based on any given criteria predetermined by the designer. Examples of these criteria are priorities, weights assigned to classes, which class pays more than others, which class should receive preferential treatment over others, etc." and Col. 8, lines 23 – 25: "When the customer starts using the bandwidth in this range, the customer is charged for such usage."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate charging an entity according to bandwidth usage, as taught by Boivie et al., in the method and system for allocating a shared resource among a plurality of devices, as in

Iverson et al. for the purpose of allocating access to a shared resource based on how much an entity pays.

**Claims 9** and **22** are rejected under 35 U.S.C. 103(a) as being unpatentable over **Iverson et al. (U.S. Patent # 6052379)** in view of **RFC 2212 ("Specification of Guaranteed Quality of Service").**

Consider **claims 9** and **22**, and as applied to claims 1 and 14, Iverson et al. disclose the claimed invention except determining a maximum latency when a bucket reaches a maximum number of credits. However, RFC 2212 clearly shows and discloses determining a maximum latency (page 3, End-to-End Behavior section: "Guaranteed service does not control the minimal or average delay of datagrams, merely the maximal queueing delay. Furthermore, to compute the maximum delay a datagram will experience, the latency of the path MUST be determined and added to the guaranteed queueing delay. (However, as noted below, a conservative bound of the latency can be computed by observing the delay experienced by any one packet)." and page 16, Guidelines for Implementors section: "However, even if latency is not advertised, this service can still be used. The simplest approach is to measure the delay experienced by the first packet (or the minimum delay of the first few packets) received and treat this delay value as an upper bound on the latency."). Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate determining a maximum latency as taught by RFC 2212 in the method and system for allocating a shared resource among a plurality of devices as in

Iverson et al. for the purpose of using the maximum latency to determine the maximum delay that communication from a given device will experience in using the shared resource.

**Claims 10** and **23** are rejected under 35 U.S.C. 103(a) as being unpatentable over **Iverson et al. (U.S. Patent # 6052379)** in view of **RFC 2212 ("Specification of Guaranteed Quality of Service")** and further in view of **Ng et al. (U.S. Patent # 6157978).**

Consider **claims 10** and **23**, and as applied to claims 9 and 22, Iverson et al. as modified by RFC 2212 disclose the claimed invention except granting immediate access to the shared resource to the device whose maximum latency has been reached. However, Ng et al. clearly show and disclose granting immediate access to the shared resource to the device which has been stalled for a time that is nearing its maximum latency (Col. 6, lines 53 – 63: "The inventors have realized that the low latency for a communications device is much shorter than the full arbitration loop when multiple multimedia agents access a shared resource such as a shared multimedia memory. Reducing the length of the loop could reduce latency, but also reduces bandwidth to the other multimedia agents. Instead, the inventors have realized that low latency can be achieved if the super-priority agent is allowed to enter the arbitration loop at any point in the loop. Then the super-priority agent does not have to wait for the beginning of the next loop." and Col. 11, lines 7 – 12: "The phantom slot following a super-priority request is used for the super-priority agent, providing a low latency. In the example of

Fig. 8, a super-priority request is received during slot A, while RT agent A is being serviced. Once agent A finishes, a phantom slot for the super-priority agent is inserted. During slot S1 (time $T_{SP}$) the super-priority agent is serviced. The latency is less than the time for slot A, $T_A$, which is much less than the entire loop time $T_{LOOP}$." and Col. 7, line 63 – Col. 8, line 6: "Fig. 4 is an arbitration loop diagram with phantom slots for use by a super-priority agent. High-priority real-time agents A, B, C, D are assigned a fixed slot in the arbitration loop. Other non-real-time (NRT) agents share a remaining NRT slot that is terminated after a timeout. Phantom slots S1, S2, S3, S4, S5 are inserted between the RT slots A, B, C, D, and the NRT slot. These phantom slots are shown with dashed outlines. Phantom slots are skipped when the super-priority agent is inactive, so that the arbitration loop collapses back to the A-B-C-D-NRT loop of the parent application."). In other words, when a super-priority agent requests access to the shared resource, it is given immediate access. Agents that have low latency requirements are designated super-priority. An obvious variation of this would be to designate a device whose bucket has reached the maximum number of credits, i.e. a device which now requires low latency, super-priority and give it immediate access to the shared resource. Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to incorporate granting immediate access to the shared resource to the device which requires low latency, as taught by Ng et al., in the method and system for allocating a shared resource among a plurality of devices as in Iverson et al. in view of RFC 2212 for the purpose of preventing underflow failures.

## *Conclusion*

Any response to this Office Action should be **faxed to** (571) 273-8300 **or mailed to**:

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

**Hand-delivered responses** should be brought to

Customer Service Window
Randolph Building
401 Dulany Street
Alexandria, VA 22314

Any inquiry concerning this communication or earlier communications from the Examiner should be directed to Leah Richmond whose telephone number is (571) 270-1774. The Examiner can normally be reached on Monday-Thursday from 9:00am to 6:00pm Eastern Standard Time.

If attempts to reach the Examiner by telephone are unsuccessful, the Examiner's supervisor, Rafael Pérez-Gutiérrez can be reached at (571) 272-7915. The fax phone number for the organization where this application or proceeding is assigned is (571) 273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you

have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free) or 703-305-3028.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist/customer service whose telephone number is (571) 272-2600.

*Leah Richmond*
L.L.R./llr

March 20, 2007